Computational Geometry 21 (2002) 139–153

Computational Geometry
Theory and Applications

www.elsevier.com/locate/comgeo

Inapproximability of finding maximum hidden sets on polygons and terrains ☆

Stephan Eidenbenz

Institute for Theoretical Computer Science, ETH Zürich, Zürich, Switzerland

Communicated by Subhash Suri; received 3 July 2000; received in revised form 20 December 2000; accepted 5 January 2001

Abstract

How many people can hide in a given terrain, without any two of them seeing each other? We are interested in finding the precise number and an optimal placement of people to be hidden, given a terrain with n vertices. In this paper, we show that this is not at all easy: The problem of placing a maximum number of hiding people is almost as hard to approximate as the MAXIMUM CLIQUE problem, i.e., it cannot be approximated by any polynomial-time algorithm with an approximation ratio of n^{ε} for some $\varepsilon > 0$, unless $\mathrm{P} = \mathrm{NP}$. This is already true for a simple polygon with holes (instead of a terrain). If we do not allow holes in the polygon, we show that there is a constant $\varepsilon > 0$ such that the problem cannot be approximated with an approximation ratio of $1 + \varepsilon$. © 2002 Published by Elsevier Science B.V.

Keywords: Hiding; Visibility; Inapproximability; Polygons; Terrains

1. Introduction and problem definition

While many of the traditional art gallery problems such as VERTEX GUARD and POINT GUARD deal with the problem of guarding a given polygon with a minimum number of guards, the problem of hiding a maximum number of objects from each other in a given polygon is intellectually appealing as well. When we let the problem instance be a terrain rather than a polygon, we obtain the following background, which is the practical motivation for the theoretical study of our problem: A real estate agency owns a large, uninhabited piece of land in a beautiful area. The agency plans to sell the land in individual pieces to people who would like to have a cabin in the wilderness, which to them means that they do not see any signs of human civilization from their cabins. Specifically, they do not want to see any other cabins. The real estate agency, in order to maximize profit, wants to sell as many pieces of land as possible.

In an abstract version of the problem we are given a terrain which represents the uninhabited piece of land that the real estate agency owns. A *terrain T* is a two-dimensional surface in three-dimensional

☆ A preliminary version of this report has appeared as [5].

E-mail address: eidenben@inf.ethz.ch (S. Eidenbenz).

PII: S0925-7721(01)00029-3

space, represented as a finite set of vertices in the plane, together with a triangulation of their planar convex hull, and a height value associated with each vertex. By a linear interpolation inbetween the vertices, this representation defines a bivariate continuous function. The corresponding surface in space is also called a 2.5-dimensional terrain. A terrain divides three-dimensional space into two subspaces, i.e., a space above and a space below the terrain, in the obvious way. In the literature, a terrain is also called a *triangulated irregular network* (TIN), see [16]. The problem now consists of finding a maximum number of lots (of comparatively small size) in the terrain, upon which three-dimensional bounding boxes can be positioned that represent the cabins such that no two points of two different bounding boxes see each other. Two points *see* each other, if the straight line segment connecting the two points does not intersect the space below the terrain. Since the bounding boxes that represent the cabins are small compared to the overall size and elevation changes in the terrain (assume that we have a mountainous terrain), we may consider these bounding boxes to be zero-dimensional, i.e., to be points on the terrain. We are now ready to formally define the first problem that we study:

Definition 1. The problem MAXIMUM HIDDEN SET ON TERRAIN asks for a set S of maximum cardinality of points on a given terrain T, such that no two points in S see each other.

In a variant of the problem, we introduce the additional restriction that these points on the terrain must be vertices of the terrain.

Definition 2. The problem MAXIMUM HIDDEN VERTEX SET ON TERRAIN asks for a set S of maximum cardinality of vertices of a given terrain T, such that no two vertices in S see each other.

In a more abstract variant of the same problem, we are given a simple polygon with or without holes instead of a terrain. A *simple polygon with holes* in the plane is given by its ordered sequence of vertices on the outer boundary, together with an ordered sequence of vertices for each hole. A *simple polygon without holes* in the plane is simply given by its ordered sequence of vertices on the outer boundary. Two points in the polygon see each other, if the straight line segment connecting the two points does not intersect the exterior (and the holes) of the polygon. Again, we can impose the additional restriction that the points to be hidden from each other must be vertices of the polygon. This yields the following four problems.

Definition 3. The problem MAXIMUM HIDDEN SET ON POLYGON WITH(OUT) HOLES asks for a set S of maximum cardinality of points in the interior or on the boundary of a given polygon P, such that no two points in S see each other.

Definition 4. The problem MAXIMUM HIDDEN VERTEX SET ON POLYGON WITH(OUT) HOLES asks for a set S of maximum cardinality of vertices of a given polygon P, such that no two vertices in S see each other.

In this paper, we propose a reduction from MAXIMUM CLIQUE to MAXIMUM HIDDEN SET ON POLYGON WITH HOLES. The same reduction with minor modifications will also work for MAXIMUM HIDDEN SET ON TERRAIN, MAXIMUM HIDDEN VERTEX SET ON POLYGON WITH HOLES, and MAXIMUM HIDDEN VERTEX SET ON TERRAIN. MAXIMUM CLIQUE cannot be approximated by a polynomial-time algorithm with a ratio of $n^{1-\varepsilon}$ unless NP = ZPP [1] and with a ratio of $n^{1/2-\varepsilon}$ unless NP = P for any $\varepsilon > 0$, where n is the number of vertices in the graph [9]. We will show that our reduction is gap-preserving (a technique proposed in [1]), and thus show inapproximability results for all four problems. MAXIMUM CLIQUE consists of finding a maximum complete subgraph of a given graph G, as usual.

For input polygons without holes, we propose a reduction from MAXIMUM 5-OCCURRENCE-2-SATISFIABILITY TO MAXIMUM HIDDEN SET ON POLYGON WITHOUT HOLES, which works for MAXIMUM HIDDEN VERTEX SET ON POLYGON WITHOUT HOLES as well.

Definition 5. The problem MAXIMUM 5-OCCURRENCE-2-SATISFIABILITY consists of finding a truth assignment for the variables of a given Boolean formula. The formula consists of disjunctive clauses with at most two literals and each variable appears in at most 5 literals. The truth assignment must satisfy a maximum number of clauses.

MAXIMUM 5-OCCURRENCE-2-SATISFIABILITY is APX-hard, which is equivalent to saying that there exists a constant $\varepsilon > 0$ such that no polynomial algorithm can achieve an approximation ratio of $1+\varepsilon$ for MAXIMUM 5-OCCURRENCE-2-SATISFIABILITY. See [3] for an introduction to the class APX and for the relationship between the two classes APX and MaxSNP, see [12] for the MaxSNP-hardness proof of MAXIMUM 5-OCCURRENCE-2-SATISFIABILITY and for the definition of MaxSNP. Please note that MaxSNP-hardness implies APX-hardness [3].

We show that our reduction is gap-preserving and thus establish the APX-hardness of MAXIMUM HIDDEN (VERTEX) SET ON POLYGON WITHOUT HOLES.

There are various problems that deal with terrains. Quite often, these problems have applications in the field of telecommunications, namely in setting up communications networks. There are some upper and lower bound results on the number of guards needed for several kinds of guards to collectively cover all of a given terrain [2]. Very few results on the computational complexity of terrain problems are known. The shortest watchtower (from where a terrain can be seen in its entirety) can be computed in time $O(n \log n)$ [17]. The problem of finding a minimum number of vertices of a terrain such that guards at these vertices see all of the terrain is NP-hard and cannot be approximated with an approximation ratio that is better than logarithmic in the number of vertices of the terrain. Similar results hold for the variation, where guards may only be placed at a certain given height above the terrain [7]. One of the most intensely studied problems on terrains is the problem of computing the shortest path between two points on the terrain. This problem can be solved in time $O(n \log^2(n))$ [10]. Many results (upper and lower bounds, as well as computational complexity results) are known for visibility problems with polygons as input structures. See [11,14,15] for an overview, as well as more recent work on the inapproximability of VERTEX/EDGE/POINT GUARD on polygons with [6] and without holes [4].

The problems MAXIMUM HIDDEN SET ON A POLYGON WITHOUT HOLES and MAXIMUM HIDDEN VERTEX SET ON A POLYGON WITHOUT HOLES are known to be NP-hard [13]. This immediately implies the NP-hardness of the corresponding problems for polygons with holes. A quite simple reduction from these polygon problems to the terrain problems (as given in Section 3) even implies the NP-hardness for the two terrain problems as well. We report the first *inapproximability results* for these problems.

[1] ZPP the class of problems that can be solved in expected polynomial time by a probabilistic algorithm that never makes an error, i.e., only the running time is stochastic.

In Section 2, we propose a reduction from MAXIMUM CLIQUE to MAXIMUM HIDDEN SET ON POLYGON WITH HOLES. We show that our reduction is gap-preserving and obtain our inapproximability results for MAXIMUM HIDDEN (VERTEX) SET ON A POLYGON WITH HOLES. We show that our proofs also work for MAXIMUM HIDDEN (VERTEX) SET ON TERRAIN with minor modifications in Section 3. In Section 4, we show the APX-hardness of MAXIMUM HIDDEN (VERTEX) SET ON POLYGON WITHOUT HOLES. We draw conclusions in Section 5.

2. Hiding in polygons with holes

We propose a reduction from MAXIMUM CLIQUE to MAXIMUM HIDDEN SET ON POLYGON WITH HOLES that constructs a polygon with holes for the hiding instance that very naturally corresponds to the input graph of the clique instance.

Suppose we are given an instance I of MAXIMUM CLIQUE, i.e., an undirected graph $G = (V, E)$, where $V = v_0, \dots, v_{n-1}$. Let $m := |E|$. Fig. 1 shows an example. We construct an instance I' of MAXIMUM HIDDEN SET ON POLYGON WITH HOLES as follows. I' consists of a polygon with holes. The polygon is basically a regular $2n$-gon with holes, but we replace every other vertex by a comb-like structure. Each hole is a small triangle designed to block the view of two combs from each other, whenever the two vertices, to which the combs correspond, are connected by an edge in the graph. Fig. 2 shows as an example the polygon with holes constructed from the graph in Fig. 1. (Note that only the solid lines are lines of the polygon and also note that the combs are not shown in Fig. 2.)

Let the regular $2n$-gon consist of vertices $v_0, v_0', \dots, v_{n-1}, v_{n-1}'$ in counterclock-wise order, to indicate that we map each vertex $v_i \in V$ in the graph to a vertex v_i in the polygon.

We need some notation, first. Let $e_{i,j}$ denote the intersection point of the line segment from v_{i-1}' to v_i' with the line segment from v_i to v_j, as indicated in Fig. 3. (Note that we make liberal use of the notation index for the vertices, i.e., v_{i+1} is strictly speaking $v_{i+1 \bmod n}$, accordingly for v_{i-1}.) Let d denote the minimum of the distances of $e_{i,j}$ from $e_{i,j+1}$, where the minimum is taken over all $i, j = 1, \dots, n-1$. Let $e_{i,j}^-$ ($e_{i,j}^+$) denote the point at distance $d/3$ from $e_{i,j}$ on the line from v_{i-1}' to v_i' that is closer to v_{i-1}' (v_i'). Let m_i be the midpoint of the line segment from vertex v_i to v_{i+1} and let m_i' be the intersection

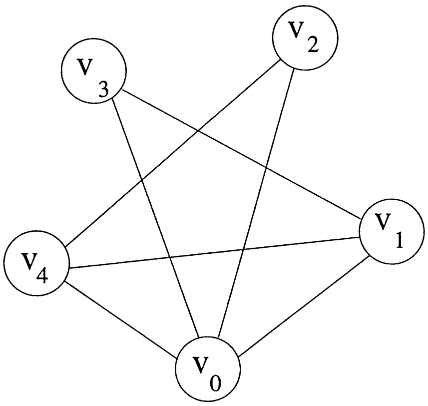

Fig. 1. A graph with five vertices.

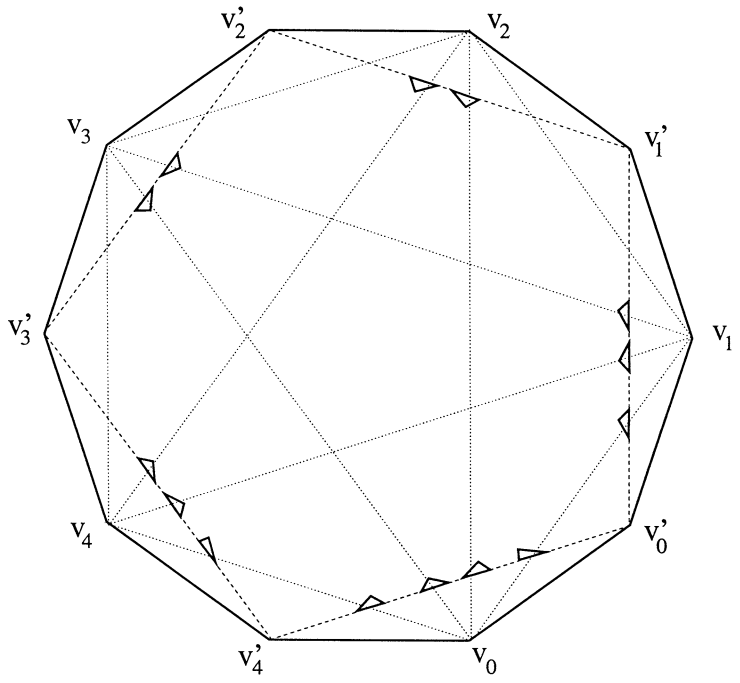

Fig. 2. Polygon constructed from the graph in Fig. 1.

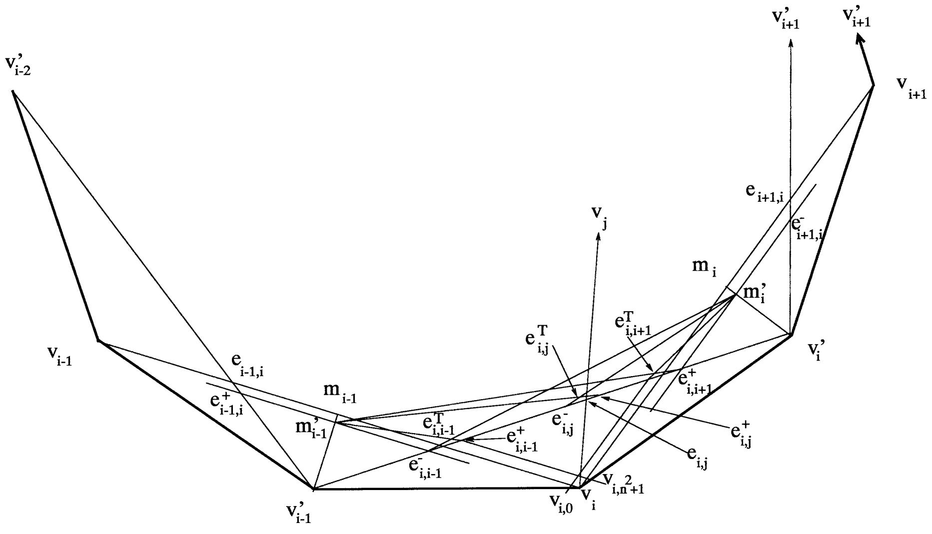

Fig. 3. Points $e^{-}_{i,j}$, $e^{+}_{i,j}$ and $e^{\mathrm{T}}_{i,j}$.

point of the line from v_i' to m_i and from $e^{+}_{i,i+1}$ to $e^{-}_{i+1,i}$ (see Fig. 3). Finally, let $e^{\mathrm{T}}_{i,j}$ denote the intersection point of the line from $e^{-}_{i,j}$ to m_i' and the line from $e^{+}_{i,j}$ to m_{i-1}'. The detailed construction of these points is shown in Fig. 3.

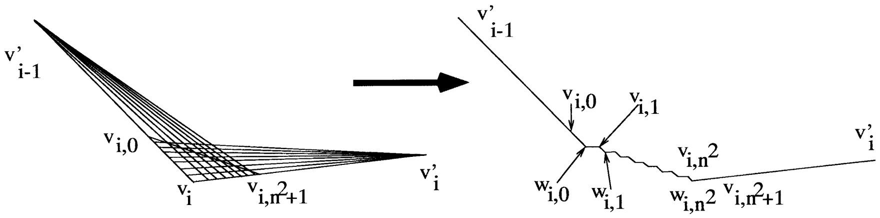

Fig. 4. Construction of the comb of v_i.

We let the triangle formed by the three vertices $e^-_{i,j}$, $e^+_{i,j}$ and $e^{\mathrm{T}}_{i,j}$ be a hole in the polygon iff there exists an edge in G from v_i to v_j. Recall Fig. 2, which gives an example. We now refine the polygon obtained so far by cutting off a small portion at each vertex v_i. For each $i \in \{0, \ldots, n-1\}$, we introduce two new vertices $v_{i,0}$ and v_{i,n^2+1} as indicated in Fig. 3. Vertex v_{i,n^2+1} is defined as the intersection point of the line that is parallel to the line from v_{i-1} to v_i and goes through point $e^+_{i,i-1}$ and of the line from v_i to v'_i. Symmetrically, vertex $v_{i,0}$ is defined as the intersection point of the line that is parallel to the line from v_{i+1} to v_i and that goes through point $e^-_{i,i+1}$ and of the line from v_i to v'_{i-1}.

We fix n^2 additional vertices $v_{i,1}, \ldots, v_{i,n^2}$ on the line segment from $v_{i,0}$ to v_{i,n^2+1} for each i as shown in Fig. 4. For a fixed i, the two vertices $v_{i,l}$ and $v_{i,l+1}$ have equal distance for all $l \in \{0, \ldots, n^2\}$. Finally, we fix n^2+1 additional vertices $w_{i,l}$ for $l \in \{0, \ldots, n^2\}$ for each i. Vertex $w_{i,l}$ is defined as the intersection point of the line from vertex v_{i-1} through $v_{i,l}$ with the line from vertex v'_i through $v'_{i,l+1}$ The polygon between two vertices v'_{i-1} and v'_i is now given by the following ordered sequence of vertices:

$$v'_{i-1}, v_{i,0}, w_{i,0}, v_{i,1}, w_{i,1}, \ldots, v_{i,n^2}, w_{i,n^2}, v_{i,n^2+1}, v'_i$$

as indicated in Fig. 4. We call the set of all triangles $v_{i,l}, w_{i,l}, v_{i,l+1}$ for a fixed i and all $l \in \{0, \ldots, n^2\}$ the *comb of* v_i.

The constructed polygon satisfies the following property.

Lemma 1. *In any feasible solution S' of the* MAXIMUM HIDDEN SET ON POLYGON WITH HOLES *instance I', at most $2n$ points in S can be placed outside the combs.*

Proof. In each of the n trapezoids $\{v'_{i-1}, v'_i, v_{i,n^2+1}, v_{i,0}\}$ (see Figs. 2 and 3), there can be at most one point, which gives n points in total. Moreover, by our construction any point p in the trapezoid $\{v'_{i-1}, v'_i, m'_i, m'_{i-1}\}$ (not in the holes) can see every point p' in the n-gon $\{v'_0, \ldots, v'_n\}$ except for points p' in any of the holes and (possibly) except for points p' in the triangles $\{v'_{i-1}, m'_{i-1}, e^+_{i-1,i}\}$ and $\{v'_i, m'_i, e^-_{i+1,i}\}$ (see Fig. 3). Therefore, all points in S' that lie in the n-gon $\{v'_0, \ldots, v'_n\}$ must lie in only one of the n polygons $\{e^+_{i-1,i}, m'_{i-1}, m'_i, e^+_{i+1,i}, v'_i, v'_{i-1}\}$. Obviously, at most n points can be hidden in any one of these polygons. □

We have the following observation, which follows directly from the construction.

Observation 1. *Any point in the comb of v_i sees the entire comb of vertex v_j, if (v_i, v_j) is not an edge in the graph. If (v_i, v_j) is an edge in the graph, then no point in the comb of v_i sees any point in the comb of v_j.*

Given a feasible solution S' of the MAXIMUM HIDDEN SET ON POLYGON WITH HOLES instance I', we obtain a feasible solution S of the MAXIMUM CLIQUE instance I as follows: A vertex $v_i \in V$ is in the solution S, iff at least one point from S' lies in the comb of v_i. To see that S is a feasible solution, assume by contradiction that it is not a feasible solution. Then, there exists a pair of vertices $v_i, v_j \in S$ with no edge between them. But then, there is by construction no hole in the polygon to block the view between the comb of v_i and the comb of v_j.We need to show that the construction of I' can be done in polynomial time and that a feasible solution can be transformed in polynomial time. There are $2n^2+1$ vertices in each of the n combs. We have additional n vertices v_i'. There are 2 holes for each edge in the graph and each hole consists of 3 vertices. Therefore, the polygon consists of $6m+2n^3+2n$ vertices. It is known in computational geometry that the coordinates of intersection points of lines with rational coefficients can be expressed with polynomial length [8]. All of the points in our construction are of this type. Therefore, the construction is polynomial. The transformation of a feasible solution can obviously be done in polynomial time.

We obtain our inapproximability result, again, by using the technique of gap-preserving reductions, which consists of transforming a promise problem into another promise problem. Let OPT denote the size of an optimum solution of the MAXIMUM CLIQUE instance I, let OPT$'$ denote the size of an optimum solution of the MAXIMUM HIDDEN SET ON POLYGON WITH HOLES instance I', let $k \leqslant n$, and let $\varepsilon > 0$.

Lemma 2. $\mathrm{OPT} \geqslant k \Rightarrow \mathrm{OPT}' \geqslant n^2 k$.

Proof. If $\mathrm{OPT} \geqslant k$, then there exists a clique in I of size k. We obtain a solution for I' of size n^2k by simply letting the n^2 vertices $w_{i,l}$ for $l \in \{0, \ldots, n^2\}$ be in the solution if and only if vertex $v_i \in V$ is in the clique. The solution thus obtained for I' is feasible (see Observation 1). □

Lemma 3. $\mathrm{OPT} < k/n^{1/2-\varepsilon} \Rightarrow \mathrm{OPT}' < n^2k/n^{1/2-\varepsilon} + 2n$.

Proof. We prove the contraposition:

$$\mathrm{OPT}' \geqslant \frac{n^2k}{n^{1/2-\varepsilon}} + 2n \quad \Longrightarrow \quad \mathrm{OPT} \geqslant \frac{k}{n^{1/2-\varepsilon}}.$$

Suppose we have a solution of I' with $n^2k/n^{1/2-\varepsilon} + 2n$ points. At most $2n$ of the points in the solution can be outside the combs, because of Lemma 1. Therefore, at least $n^2k/n^{1/2-\varepsilon}$ points must be in the combs. From the construction of the combs, it is clear that at most n^2 points can hide in each comb. Therefore, the number of combs that contain at least one point from the solution is at least

$$\frac{n^2k/n^{1/2-\varepsilon}}{n^2} = \frac{k}{n^{1/2-\varepsilon}}.$$

The transformation of a solution as described above yields a solution of I with at least $k/n^{1/2-\varepsilon}$ vertices. □

Lemmas 2 and 3 transform the NP-hard promise problem of MAXIMUM CLIQUE, where we are promised that the optimum solution consists of either at least k or strictly less than $k/n^{1/2-\varepsilon}$ vertices, into the NP-hard promise problem of MAXIMUM HIDDEN SET ON POLYGON WITH HOLES, where we are promised that an optimum solution consists of either at least n^2k or strictly less than $n^2k/n^{1/2-\varepsilon} + 2n$

hidden points. Thus, MAXIMUM HIDDEN SET ON POLYGON WITH HOLES cannot be approximated with an approximation ratio of

$$\frac{n^2k}{n^2k/n^{1/2-\varepsilon}+2n} > \frac{n^{1/2-\varepsilon}}{2},$$

where we have assumed that $k \geqslant 2$. We need to express the number of graph vertices n by the number of polygon vertices $|I'|$ of the polygon of instance I'. Note that $|I'| \leqslant 10n^3$ and therefore $n \geqslant |I'|^{1/3}/3$. Thus,

$$\frac{n^{1/2-\varepsilon}}{2} \geqslant \frac{|I'|^{1/6-\varepsilon/3}/3^{1/2-\varepsilon}}{2} > \frac{|I'|^{1/6-\varepsilon/3}}{4}.$$

This yields the main result of this section.

Theorem 1. MAXIMUM HIDDEN VERTEX SET ON POLYGON WITH HOLES *cannot be approximated by any polynomial time algorithm with an approximation ratio of* $|I'|^{1/6-\gamma}/4$*, where* $|I'|$ *is the number of vertices in the polygon, and where* $\gamma > 0$*, unless* NP $=$ P.

If we restrict the hidden set to contain only vertices, we can use the same construction. Actually, we do not need the combs, as our construction guarantees that in any solution there can be at most 2 points hiding at vertices other than v_i. This leads to a different promise problem of MAXIMUM HIDDEN VERTEX SET OF POLYGON WITH HOLES that the promise problem of MAXIMUM CLIQUE is mapped to, namely the promise problem, where we are promised that an optimum solution consists of either at least k vertices or strictly less than $k/n^{1/2-\varepsilon}+2$ vertices. Straightforward analysis, using the fact that $|I'| \leqslant 5n^2$, leads to the following result.

Theorem 2. MAXIMUM HIDDEN VERTEX SET ON POLYGON WITH HOLES *cannot be approximated by any polynomial time algorithm with an approximation ratio of* $|I'|^{1/4-\gamma}/4$*, where* $|I'|$ *is the number of vertices in the polygon, and where* $\gamma > 0$*, unless* NP $=$ P.

3. Hiding in terrains

Theorem 3. *The problems* MAXIMUM HIDDEN SET ON TERRAIN (MAXIMUM HIDDEN VERTEX SET ON TERRAIN) *cannot be approximated by any polynomial time algorithm with an approximation ratio of* $(|I''|^{1/6-\gamma}/4)(|I''|^{1/4-\gamma}/4)$*, where* $|I''|$ *is the number of vertices in the terrain, and where* $\gamma > 0$*, unless* NP $=$ P.

Proof. The proof very closely follows the lines of the proof for the inapproximability of MAXIMUM HIDDEN (VERTEX) SET ON POLYGON WITH HOLES. We use the same construction, but given the polygon with holes of instance I' we create a terrain (i.e., instance I'') by simply letting all the area outside the polygon (including the holes) have height h for some constant $h > 0$ and by letting the area in the interior have height 0.

We add four vertices to the terrain by introducing a rectangular bounding box around the regular $2n$-gon. This yields a terrain with vertical walls, which can be easily modified to have steep but not vertical walls, as required by the definition of a terrain. Finally, we triangulate the terrain. □

4. Hiding in polygons without holes

We propose a gap-preserving reduction from MAXIMUM 5-OCCURRENCES-2-SATISFIABILITY to MAXIMUM HIDDEN SET ON POLYGON WITHOUT HOLES, which allows us to prove the APX-hardness of MAXIMUM HIDDEN SET ON POLYGON WITHOUT HOLES. The same reduction will also work for MAXIMUM HIDDEN VERTEX SET ON POLYGON WITHOUT HOLES with minor modifications.

Suppose we are given an instance I of MAXIMUM 5-OCCURRENCE-2-SATISFIABILITY, which consists of n variables $x_0, \ldots, x_{n-1}$ and m clauses $c_0, \ldots, c_{m-1}$. We construct a polygon without holes, i.e., an instance I' of MAXIMUM HIDDEN SET ON POLYGON WITHOUT HOLES, which consists of clause patterns and variable patterns, as shown schematically in Fig. 5. The construction consists of the following building blocks:

We have a variable pattern for each variable that consists of a TRUE- and a FALSE-leg. We have a clause pattern for each clause that is a simple zig-zag line forming 3 dents. The variable patterns contain a small dent, which we will call "cone", for each occurrence of the variable in the input satisfiability formula. These cones "connect" the variable patterns with the clause patterns.

We construct a variable pattern for each variable x_i as indicated in Fig. 6. Each variable pattern consists of a TRUE- and a FALSE-leg. Each leg has on its left boundary a maximum of five triangle-shaped dents with vertices f_k for $k = 1, \ldots, 5$. Each of these dents represents the lower part of a cone that connects the variable pattern to a clause pattern, in which the variable appears as a literal, as indicated in Fig. 7. Since these dents are triangle-shaped, we call the whole dent the *triangle of* f_k. All these triangles are attached to a single line as indicated in Fig. 6 on their right side, i.e., the "right" line segments of each triangle in a leg are collinear.

Each variable pattern contains at its right side exactly four dents with vertices v_1, v_2, v_3 and w. The construction is such that a point that is hiding in the triangle of any f_k sees the triangles of v_1, v_2 and v_3 completely, but it does not see the triangle of w. Therefore, we can hide in each leg a point at vertex w and additional points either in the triangles of f_k on the left side or in the triangles of v_k on the right side, but not in both. The idea is the following: if the variable is TRUE, then we hide points in the TRUE-leg

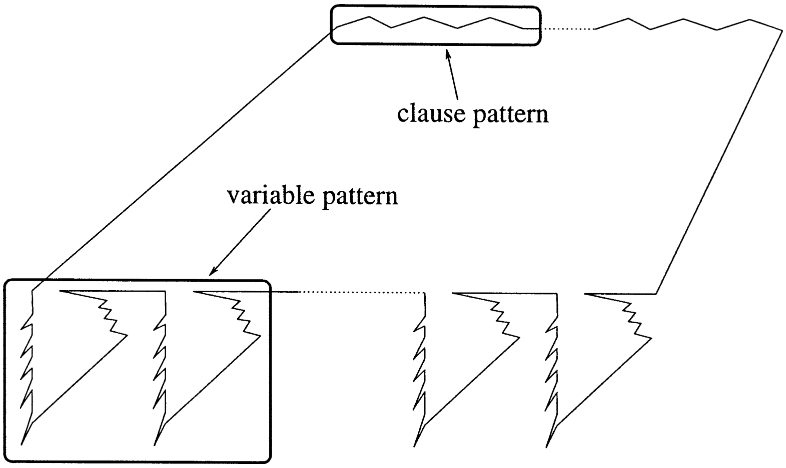

Fig. 5. Schematic construction.

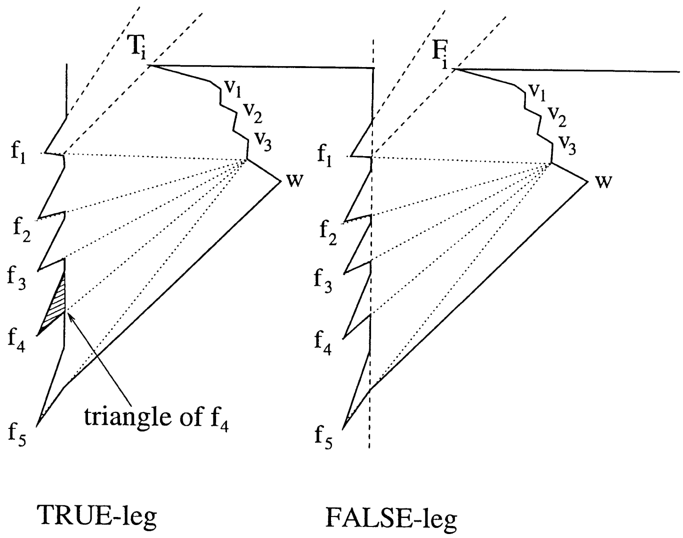

Fig. 6. Variable pattern.

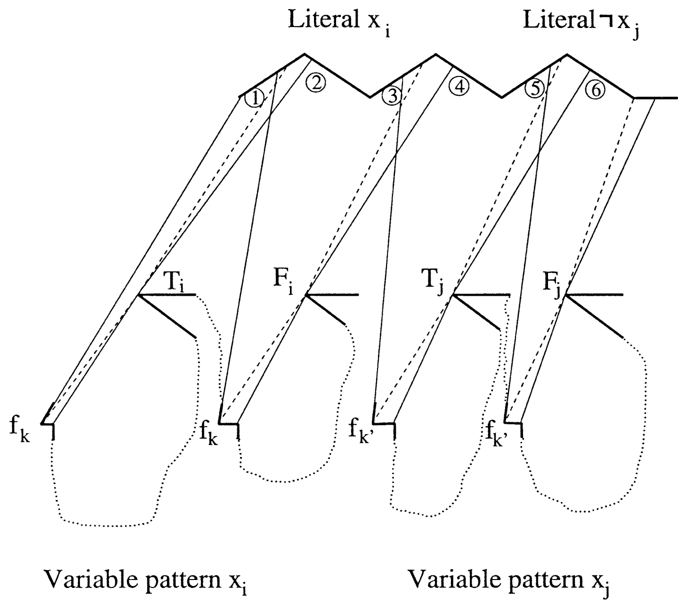

Fig. 7. Clause pattern with cones.

of the variable pattern at vertices f_k, for $k = 1, \ldots, 5$ and at vertex w. In the FALSE-leg, we hide points at vertices v_1, v_2, v_3 and w.

For each clause c_i we construct a clause pattern as indicated in Fig. 7. A clause pattern consists of three dents, where the left and the right dent represent the literals of the clause. The middle dent represents the truth value of the clause. The construction is such that we can hide three points in the clausepattern (i.e., one in each dent), exactly if the clause is satisfied. We can hide only two points (one in the left and one the right dent), otherwise. To achieve this we connect the variable patterns to the clause patterns with cones as illustrated in Fig. 7 for two variables x_i and x_j and a clause $(x_i, \neg x_j)$. This works accordingly for other types of clauses. Cones are drawn as thin solid lines. They are not part of the polygon boundary, but merely help in the construction.

We will show that this reduction is gap-preserving, i.e., it maps an NP-hard promise problem of MAXIMUM 5-OCCURRENCE-2-SATISFIABILITY to an NP-hard promise problem of MAXIMUM HIDDEN SET ON POLYGON WITHOUT HOLES. The reduction has the following properties:

Lemma 4. *If there exists a truth assignment S to the variables of the* MAXIMUM 5-OCCURRENCE-2-SATISFIABILITY *instance I that satisfies at least $(1-\varepsilon)m$ clauses, then there exists a solution S' of the* MAXIMUM HIDDEN SET WITHOUT HOLES *instance I' with $|S'| \geqslant 10n + 2m + (1-\varepsilon)m$.*

Proof. If variable x_i is TRUE in S, then we let the vertices $f_1, \ldots, f_5$ and w of the TRUE-leg of x_i, as well as the vertices v_1, v_2, v_3 and w of the FALSE-leg of x_i be in the solution S'; vice-versa if x_i is FALSE in S. This gives us $10n$ points in S'.

The remaining points for S' are in the clause patterns. Fig. 7 shows the clause pattern for a clause $(x_i, \neg x_j)$ [2], together with the cones that link the clause pattern to the corresponding variable patterns. Remember that these cones are not part of the polygon boundary. To understand Fig. 7, assume x_i is assigned the value FALSE and x_j is assigned the value TRUE, i.e., the clause $(x_i, \neg x_j)$ is not satisfied. Then there is a point in the solution that sits at vertex f_k (for some k) in the FALSE-leg of x_i and a point that sits at vertex $f_{k'}$ (for some k') in the TRUE-leg of x_j. In this case, we can have only *two* additional points in the solution S' at points ①, ⑥. In the remaining three cases, where the variables x_i and x_j are assigned truth values such that the clause is satisfied, we can have *three* additional points in S' at ①–⑥: If x_i and x_j are both FALSE in the solution S, then we hide points in the clause pattern at points ①, ④ and ⑤. If x_i and x_j are both TRUE, then we hide points in the clause pattern at points ②, ③ and ⑥. If x_i is TRUE and x_j is FALSE, then we can hide points in the clause pattern at point ② and ⑤, and one additional point at either ③ or ④.

Therefore, we have 2 points from all unsatisfied clauses and 3 points from all satisfied clauses, i.e., $2\varepsilon m + 3(1-\varepsilon)m$ points that are hidden in the clause patterns. Thus, $|S'| \geqslant 10n + 2m + (1-\varepsilon)m$, as claimed. □

Lemma 5. *If there exists a solution S' of I' with $|S'| \geqslant 10n + 3m - (\varepsilon + \gamma)m$, then there exists a variable assignment S of I that satisfies at least $(1-\varepsilon-\gamma)m$ clauses.*

Proof. For any solution S', we can assume that in each leg of each variable pattern, all points in S' are either in the triangles of vertices $f_1, \ldots, f_5$ and w, or in the triangles of vertices v_1, v_2, v_3 and w. To see this, note that there can be at most one point in each leg outside the triangles. This point either sees w or at least one of the triangles v_1, v_2, v_3. In the first case, we can move the point to w. In the second case, we move it to the v_l in sight. This results again in a feasible solution. Furthermore, any point in any triangle of $f_1, \ldots, f_5$ sees the triangles of v_1, v_2, v_3 completely (and vice-versa). [3]

---

[2] The proofs work accordingly for other types of clauses, such as (x_i, x_j).

[3] We need the triangle of w to ensure that, if there are points hiding in the triangles of the f_k's, there can be no additional point hiding outside the triangles in the leg.

As a first step in transforming the solution S', we move the points that are hiding in the triangles to the vertices f_k, v_1, v_2, v_3 and w, respectively. This transformation is obviously no problem for the triangles of v_1, v_2, v_3 and w. Moving the points in the triangles of vertices f_k will slightly change the cones that they see as indicated in Fig. 7 through dashed lines. However, we can still position at least the same number of points in the clause patterns. Thus, this transformation yields a feasible hidden set of points that consists of at least as many points as the original solution.

We now describe how to transform the solution S' (with $|S'| \geqslant 10n + 3m - (\varepsilon + \gamma)m$) in such a way that it remains feasible, that its size (i.e., the number of hidden points) does not decrease, and that we will be able to assign truth values to the variables. Hence our goal is that at the end, we have for each variable pattern the six points at $f_1, \ldots, f_5$, and w from one leg in the solution and the 4 points v_1, v_2, v_3 and w from the other leg. Thus, we can easily obtain a truth assignment for the variables by letting variable x_i be TRUE iff the six points at $f_1, \ldots, f_5$ and w from the TRUE-leg are in the solution. We now show how to transform a feasible hidden set of points into one that obeys to these additional properties without decreasing the number of hidden points.

Hence, for any variable x_i we show how to transform the two legs corresponding to x_i into the desired configuration. Note that the construction of the TRUE-leg and the FALSE-leg are symmetric. For the case analysis, we may therefore assume without loss of generality that there are at least as many points hiding in the f_k triangles of the TRUE-leg as in the f_k triangles of the FALSE-leg.

- If there are 5 points hidden at vertices f_k of the TRUE-leg of x_i and 5, 4, 3, 2 or 1 point(s) at vertices f_k of the FALSE-leg, then we delete the points in the FALSE-leg and set 3 hidden points at vertices v_l of the FALSE-leg. This yields a better solution, since the difference of guards in the variable pattern is -2, -1, 0, $+1$, $+2$ and we can position 5, 4, 3, 2, 1 additional guards in the dents of the clause patterns that correspond to literals of x_i. These additional guards could not have been placed in the original solution, since the whole area of the dents was seen by the points at vertices f_k of both the TRUE- and the FALSE-leg.
- If there are 4 points hidden at vertices f_k of the TRUE-leg of x_i and 4, 3, 2 or 1 point(s) at vertices f_k of the FALSE-leg, then we delete the points in the FALSE-leg and set 3 hidden points at vertices v_l of the FALSE-leg. Moreover, we place an additional point at the one vertex f_k in the TRUE-leg, where there is not already a point hiding. Again, this yields a better solution, since the difference of guards in the variable pattern is 0, $+1$, $+2$, $+3$ and we can position at least 3, 2, 1, 0 additional guards in the dents of the clause patterns that correspond to literals of x_i, which was impossible before. Because of the additional guard at f_k in the TRUE-leg, we might lose at most one hidden point in a middle dent of a clause pattern.
- If we have 4 points hidden at vertices f_k of the TRUE-leg and no points at vertices f_k of the FALSE-leg, then we place 1 additional guard at the one vertex where there is not already a point hidden. This yields an equally good solution, since we have 1 additional points in the variable pattern and at most 1 fewer point in the middle dents of clause patterns.
- If there are at most 3 points hidden at the vertices f_k or the TRUE-leg, we first remove all the points in the f_k triangles of the TRUE- and the FALSE-leg and put points into the v_k triangles of both legs instead. This does not reduce the number of hiding points. We then argue as follows. Let us assume without loss of generality that the literal x_i appears more often in the formula than $\neg x_i$. In this case, we remove the 3 hiding points in the v_l triangles of the TRUE-leg and replace them by 5 points in the f_k triangles of the TRUE-leg. This yields a better or equally good solution, since we have 2 additional

guards in the variable pattern and we lose at most 2 guards in the middle dents of clause patterns as we falsify at most 2 clauses.

As a last step in the transformation, we add points at each vertex w in each leg, if there are no points hiding there already.

Thus, the transformed solution S' consists of at least $10n + 3m - (\varepsilon + \gamma)m$ points, $10n$ of which lie in the variable patterns. At most 3 points can lie in each clause pattern. If 3 points lie in a clause pattern, then this clause is satisfied. Therefore, if 2 points lie in each clause pattern, there are still at least $(1 - \varepsilon - \gamma)m$ additional points in S'. These must lie in clause patterns as well. Therefore, at least $(1 - \varepsilon - \gamma)m$ clauses are satisfied. □

Lemmas 4 and 5 transform the promise problem of MAXIMUM 5-OCCURRENCE-3-SATISFIABILITY into a promise problem of MAXIMUM HIDDEN SET WITHOUT HOLES. In the promise problem of MAXIMUM 5-OCCURRENCE-2-SATISFIABILITY, we are promised that an optimum solution either satisfies at least $(1-\varepsilon)m$ clauses or strictly less than $(1-\varepsilon-\gamma)m$ clauses for some constant $\varepsilon, \gamma > 0$. For small enough values [4] of $\varepsilon, \gamma > 0$, it is NP-hard to decide, which of the two cases is true. This follows from the fact that MAXIMUM 5-OCCURRENCE-3-SATISFIABILITY is APX-complete. In the promise problem of MAXIMUM HIDDEN SET ON POLYGON WITHOUT HOLES, we are promised that either at least $10n + 3m - \varepsilon m$ points can be hidden from each other, or strictly less than $10n + 3m - (\varepsilon + \gamma)m$ points can be hidden from each other. Again, it is NP-hard to decide, what is true. Thus, MAXIMUM HIDDEN SET ON POLYGON WITHOUT HOLES cannot be approximated with an approximation ratio of

$$\begin{aligned}\frac{10n + 3m - \varepsilon m}{10n + 3m - (\varepsilon + \gamma)m} &= \frac{10n + 3m - \varepsilon m - \gamma m}{10n + 3m - \varepsilon m - \gamma m} + \frac{\gamma m}{10n + 3m - \varepsilon m - \gamma m} \\ &= 1 + \frac{\gamma m}{10n + 3m - \varepsilon m - \gamma m} \geqslant 1 + \frac{\gamma m}{m(33 - \varepsilon - \gamma)} \geqslant 1 + \frac{\gamma}{33}.\end{aligned}$$

We have used that $m \geqslant n/3$. Thus:

Theorem 4. MAXIMUM HIDDEN SET ON POLYGON WITHOUT HOLES *is APX-hard.*

If we restrict the hidden set to consist of only vertices, we can use the same reduction and the same analysis with the modification that we introduce additional vertices in each clause pattern. More specifically, we replace each edge of all dents of the clause patterns by two, slightly convex edges that have their common endpoint right where the corresponding point ①–⑥ from Fig. 7 is. Thus, the result carries over.

Theorem 5. MAXIMUM HIDDEN VERTEX SET ON POLYGON WITHOUT HOLES *is APX-hard.*

5. Conclusion

We have shown that the problems MAXIMUM HIDDEN (VERTEX) SET ON POLYGON WITH HOLES and MAXIMUM HIDDEN (VERTEX) SET ON TERRAIN are almost as hard to approximate as MAXIMUM

[4] We need ε as a second parameter, since we can find out in polynomial time, whether all clauses of a MAXIMUM 5-OCCURRENCE-2-SATISFIABILITY instance are satisfiable [12]; only the optimization version of MAXIMUM 5-OCCURRENCE-2-SATISFIABILITY is NP-hard.

CLIQUE. We could prove for all these problems an inapproximability ratio of $O(|I'|^{1/3-\gamma})$, but under the assumption that coR $\neq$ NP, using the stronger inapproximability result for MAXIMUM CLIQUE from [9]. Furthermore, we have shown that MAXIMUM HIDDEN (VERTEX) SET ON POLYGON WITHOUT HOLES is APX-hard. As for approximation algorithms, an approximation algorithm for all considered problems that simply returns a single vertex achieves an approximation ratio of n. No approximation algorithms are known that achieve approximation ratios of $o(n)$.

Note that our proofs can easily be modified to work as well for polygons or terrains, where no three vertices are allowed to be collinear.

We have classified the problems MAXIMUM HIDDEN (VERTEX) SET ON POLYGON WITH HOLES and MAXIMUM HIDDEN (VERTEX) SET ON TERRAIN to belong to the class of problems inapproximable with an approximation ratio of n^{ε} for some $\varepsilon > 0$, as defined in [1]. The APX-hardness results for the problems for polygons without holes, however, do not precisely characterize the approximability characteristics of these problem. The gap between the best (known) achievable approximation ratio (which is n) and the best inapproximability ratio is still very large for these problems and should be closed in future research.

References

[1] S. Arora, C. Lund, Hardness of approximations, in: D. Hochbaum (Ed.), Approximation Algorithms for NP-Hard Problems, PWS Publishing Company, 1996, pp. 399–446.
[2] P. Bose, T. Shermer, G. Toussaint, B. Zhu, Guarding polyhedral terrains, Computational Geometry 7 (1997) 173–185.
[3] P. Crescenzi, V. Kann, A compendium of NP optimization problems, in: G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, M. Protasi (Eds.), Complexity and Approximation. Combinatorial Optimization Problems and Their Approximability Properties, Springer, Berlin, 1999. Also available at: http://www.nada.kth.se/theory/compendium/compendium.html.
[4] S. Eidenbenz, Inapproximability results for guarding polygons without holes, in: Proc. of ISAAC ’98, Lecture Notes in Computer Science, Vol. 1533, 1998, pp. 427–436.
[5] S. Eidenbenz, How many people can hide in a terrain?, in: Proc. of ISAAC ’99, Lecture Notes in Computer Science, Vol. 1741, 1999, pp. 184–194.
[6] S. Eidenbenz, C. Stamm, P. Widmayer, Inapproximability of some Art Gallery problems, in: Proc. of 10th Canadian Conf. Computational Geometry, 1998, pp. 64–65.
[7] S. Eidenbenz, C. Stamm, P. Widmayer, Positioning guards at fixed height above a terrain – an optimum inapproximability result, in: Proc. of ESA ’98, Lecture Notes in Computer Science, Vol. 1461, 1998, pp. 187–198.
[8] J.E. Goodman, J. O’Rourke (Eds.), Handbook of Discrete and Computational Geometry, CRC Press, Boca Raton, FL, 1997.
[9] J. Hastad, Clique is hard to approximate within $n^{1-\varepsilon}$, in: Proc. of the Symposium on Foundations of Computer Science, 1996.
[10] S. Kapoor, Efficient computation of geodesic shortest paths, in: Proc. of the Symposium on Theory of Computing, 1999, pp. 770–779.
[11] J. O’Rourke, Art Gallery Theorems and Algorithms, Oxford University Press, New York, 1987.
[12] C.H. Papadimitriou, Computational Complexity, Addison-Wesley, Reading, MA, 1994.
[13] T. Shermer, Hiding people in polygons, Computing 42 (1989) 109–131.
[14] T. Shermer, Recent results in Art Galleries, Proc. of the IEEE (1992).

[15] J. Urrutia, Art Gallery and illumination problems, in: J.-R. Sack, J. Urrutia (Eds.), Handbook on Computational Geometry, 1998.
[16] M. van Kreveld, Digital elevation models and TIN algorithms, in: M. van Kreveld et al. (Eds.), Algorithmic Foundations of Geographic Information Systems, Lecture Notes in Computer Science, Vol. 1340, Springer, New York, 1997, pp. 37–78.
[17] B. Zhu, Computing the shortest watchtower of a polyhedral terrain in $O(n \log n)$ time, Computational Geometry 8 (1997) 181–193.